Submitted to the 5th USA/Europe Air Traffic Management R&D Seminar, Budapest, Hungary; June 2003

METHODOLOGY FOR GENERATING CONFLICT SCENARIOS BY TIME SHIFTING RECORDED TRAFFIC DATA

Mike M. Paglione★
FAA William J. Hughes Technical Center
Atlantic City International Airport, NJ 08405, U.S.A.

Robert D. Oaks◆
Veridian IT Services, Inc.
Egg Harbor Township, NJ 08234, U.S.A.

Karl D. Bilimoria✪
NASA Ames Research Center
Moffett Field, CA 94035, U.S.A.

Abstract

A methodology is presented for generating conflict scenarios that can be used as test cases to estimate the operational performance of a conflict probe. Recorded air traffic data is used to preserve real-world errors that affect the performance of conflict probes. However, due to controller actions to separate traffic, such data generally does not contain actual violations of legal separation standards. Therefore, the track data is time shifted to create traffic scenarios featuring conflicts with characteristic properties similar to those encountered in typical air traffic operations. First, a reference set of conflicts is obtained from trajectories that are computed using birth points and nominal flight plans extracted from recorded traffic data. Distributions are obtained for several primary properties (e.g., encounter angle) that are most likely to affect the performance of a conflict probe. A genetic algorithm is then utilized to determine the values of time shifts for the recorded track data so that the primary properties of conflicts generated by the time shifted data match those of the reference set. This methodology is successfully demonstrated using recorded traffic data for the Memphis Air Route Traffic Control Center; a key result is that the required time shifts are less than 5 min for 99% of the tracks. It is also observed that close matching of the primary properties used in this study additionally provides a good match for some other secondary properties.

★Conflict Probe Assessment Team Lead, Simulation and Analysis Branch, ACB-330; *mike.paglione@faa.gov*
◆Senior Computer Systems Analyst; *rdoaks@acm.org*
✪Research Scientist, Automation Concepts Research Branch, Mail Stop 210-10; *Karl.D.Bilimoria@nasa.gov*

Introduction

A conflict is a predicted violation of minimum separation standards. A conflict probe is an air traffic management decision support tool that can detect conflicts, using information on aircraft position, speed, and flight plans, along with forecasts of wind and temperature profiles. Various approaches to conflict detection have been proposed; a survey of these methods is presented in [1]. A conflict probe would be especially useful in a future Free Flight environment [2], which is expected to have a less structured traffic pattern compared to the current operating environment.

A complete evaluation of a conflict probe has two complementary aspects: qualitative and quantitative. A qualitative evaluation generally involves real-time testing of conflict probe features and user interface through human-in-the-loop simulations and field tests; for example, [3], [4], and [5] describe real-time testing of various conflict probe capabilities. A quantitative evaluation generally involves non-real-time testing directed at the conflict prediction "engine" that underlies the features and user interface of a conflict probe. A comprehensive methodology for quantitative evaluation of a conflict probe is presented in [6]; an application of this evaluation methodology has been reported in [7]. Generic metrics for quantitative evaluation are available in [8]. Conflict probe performance metrics are presented in [3], using a hybrid approach involving data collection and transformation models applied to a recorded air traffic scenario.

Quantitative evaluation of a conflict probe requires a test scenario containing conflicts similar in nature to those encountered in typical air traffic operations. One possible approach would be to use synthetic track data generated in a simulation environment that attempts to model aircraft trajectories in the presence of various error sources. However, conflict probe performance degradation is primarily a manifestation of real-world effects that are difficult to model accurately, e.g., flight intent errors, wind model errors, aircraft dynamics modeling errors, aero-propulsive modeling errors, navigation errors, and velocity (speed and heading) errors due to radar tracker noise. For example, a statistical analysis of the influence of speed errors on conflict probe performance is presented in [9].

The objective of this work is to develop conflict scenarios using recorded air traffic data in order to preserve real-world effects. This is a challenging task because real traffic data includes the effects of controller actions to separate traffic; hence operational conflicts (actual losses of separation) are generally not present in such data. In [6], the conflict parameters were appropriately expanded for the purposes of the evaluation, and the conflict probe was evaluated on its ability to detect the corresponding "pseudo conflicts" in the recorded track data. In this work, the conflict parameters are maintained at their operational values, but the recorded track data is time-shifted to create conflicts with properties similar those encountered in actual air traffic operations.

The details of the time shifting methodology are presented in the next section, while the section following outlines the implementation of a genetic algorithm utilized for determining the time shifts in the recorded track data. The time shifting methodology is then successfully demonstrated using recorded air traffic data from the Memphis Air Route Traffic Control Center (ARTCC).

Time Shifting Methodology

The operational separation criteria for U.S. en route flight correspond to a horizontal separation standard of 5 nm and a vertical separation standard of 1,000 ft (2,000 ft if neither aircraft is below FL290). An operational conflict (actual loss of separation) occurs when these separation criteria are violated, e.g., two aircraft flying at FL250 are separated horizontally by less than 5 nm. Since real traffic data includes the effects of air traffic controller actions to maintain separation, operational conflicts do not generally exist in such a data set. Therefore a conflict probe cannot be evaluated using unmodified field data and an operational conflict window. In order to utilize actual traffic data for conflict probe evaluation, some adjustments must be made. However, these adjustments cannot be made arbitrarily.

Primary Properties of a Conflict Set

Based on the authors' prior experience with conflict probe evaluations (e.g., [6]), it is known that the performance of a conflict probe (as measured by missed/false alert rates) is strongly influenced by the characteristic properties of the conflicts themselves. For example, it is relatively easy for a conflict probe to correctly predict an opposing (encounter angle near 180 deg) collision conflict (zero distance at closest approach) between two cruising aircraft. Conversely, it is relatively difficult to correctly predict a trailing (encounter angle near 0 deg) grazing conflict (separation just below the minimum standard) between a climbing aircraft and a descending aircraft. Hence a conflict probe will perform poorly if evaluated with a traffic scenario that contains a large percentage of "difficult" conflicts.

For the purposes of this work, the primary properties of a conflict set are the following: (1) number of conflicts, and the distributions of: (2) encounter angle, (3) minimum horizontal separation, (4) minimum vertical separation, and (5) vertical flight phase (level or transitioning) of aircraft at conflict start.

Determination of Primary Properties

The purpose of conflict probe performance testing is to estimate how the conflict probe would behave under actual operational conditions. Hence the conflict scenario used for the performance evaluation should reflect the properties of conflicts encountered in actual operations. However, as stated earlier, it is not possible to determine conflict properties using recorded traffic data, because of the effects of controller actions to separate traffic. It may be theoretically possible to reconstruct what would have happened if the controllers had not intervened (e.g., analyzing voice tapes, debriefing controllers, having an observer sit next to each controller), but it is impractical to do this frequently on a large scale.

The next best option is to determine conflict property distributions from a set of trajectories generated by a high-fidelity simulation that utilizes birth points (initial conditions) and nominal flight plans extracted from recorded traffic data. For example, an aircraft's 3-D position report and active flight plan at hand-off to an ARTCC can be used to generate a trajectory through that ARTCC. This synthesized trajectory is an approximation of the actual trajectory that would have resulted in the hypothetical situation where there are no controller actions after hand-off to the ARTCC. Analysis of such trajectories for several hundred aircraft (generated over a time interval of a few hours) will yield a reference set of conflicts whose primary properties can be determined.

It is noted that these simulated trajectories will not accurately reflect real-world error sources, and are not intended to be used for conflict probe evaluation. They serve only the purpose of providing a reference set of conflicts whose primary properties are extracted for later use, as described below.

Time Shifting

The recorded tracks preserve real-world error sources, but they generally do not contain any conflicts. Therefore a time shifting process is employed to move each flight forward or backward in time by a small amount (a few minutes), in such a way that the time shifted tracks contain conflicts whose primary properties closely match those of the reference set described above. This time shifted track data can then be used as a traffic scenario for conflict probe evaluation. A schematic of the scenario generation process is presented in Fig. 1.

It is noted that time shifting a specific track simply changes (by the same amount) the time stamps associated with each 3-D position report along that track. Small time shifts are desirable so that the conflict probe evaluation can be conducted with data that is substantially similar to the recorded traffic data.

There are various techniques that could be used to determine a set of small time shifts that will satisfy the constraints of replicating the primary properties of the reference conflict set. For this work, a genetic algorithm was utilized to determine the values of time shifts, as described in the following section.

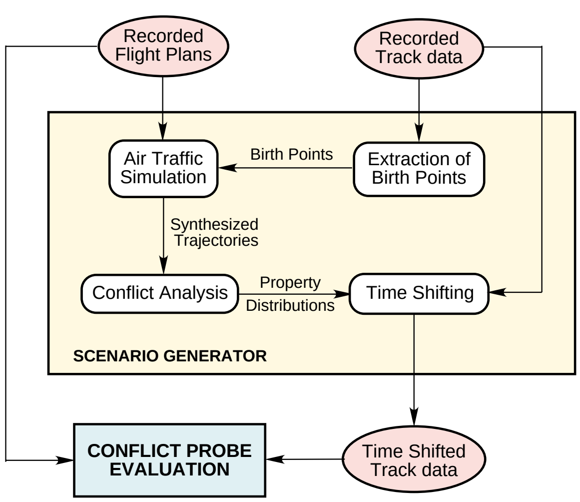

Figure 1. Schematic of Scenario Generator

Genetic Algorithm Implementation

Genetic algorithms derive their behavior from an analogy to the processes of biological evolution. They utilize a population of “chromosomes” that encode potential solutions to the problem, a fitness function that assigns a score to each potential solution, selection of a parent population according to a fitness criterion, crossover to create an offspring population, and mutation that randomly introduces new solutions into the population. A detailed treatment of genetic algorithms can be found in [10].

Genetic algorithms have been applied to a number of air traffic management problems. For example, [11] and [12] present studies in which a genetic algorithm was used for sector assignment. [13] and [14] describe studies in which a genetic algorithm was used to reduce air congestion. [15] and [16] present studies in which a genetic algorithm was used for surface management. [17] and [18] describe studies in which a genetic algorithm was used for conflict detection and resolution.

For this study, the objective is to time shift the recorded tracks so that certain properties of the resulting conflict set match those of the reference set. Using the terminology of genetic algorithms, the time shift of each flight is a “gene” on a chromosome that represents a vector of N time shifts (for a set of N flights). Hence each chromosome is a potential solution to the problem at hand.

Chromosomes are evaluated using a fitness function. A detailed description of the fitness function used in this work is given in [19]; for the purposes of this paper it is sufficient to state that the fitness function provides a value between 0.0 and 1.0, where a score of 1.0 corresponds to a chromosome that meets all imposed constraints within specified tolerances. Each constraint is a requirement to match a “slice” of a primary conflict property; e.g., 42 (±4) of the conflicts must have encounter angles between 30 and 60 deg. In this work, there are five primary properties with varying numbers of slices (bins), resulting in a total of 20 constraints. Hence obtaining a chromosome with a fitness score of 1.0 means that a set of time shifted tracks has been found whose primary property distributions closely match those of the reference set.

An initial population of 20 chromosomes, corresponding to 20 sets of initial guesses, was constructed from a normal (Gaussian) distribution of numbers with zero mean and a standard deviation of 100 sec. This initial set of 20 chromosomes represents the first generation of solutions.

A genetic algorithm uses an evolutionary process to determine successive generations of chromosomes. The evolutionary process has three steps: parent selection, crossover, and mutation. The parent selection process selects pairs of chromosomes as "parents" for the next generation's population based on each chromosome's fitness function. A *sigma scaling selection* technique was used, which favors chromosomes with a fitness value close to the average fitness of the current population. Once the parents have been selected, the crossover process randomly swaps a certain number of genes between each pair of parents. A *two-point crossover* technique was utilized, in which two loci points were randomly selected and the genes lying between these two points were exchanged across the two parent chromosomes. This process was conducted with a probability of crossover set to 0.75. After crossover, the genetic algorithm initiates the mutation process in which some genes are randomly changed to another value. This process was conducted with a probability of mutation set to 0.01. An *elitism* technique was also used, implemented as follows. The best (highest fitness score) four chromosomes were retained prior to the parent selection step. After completion of the parent selection, crossover, and mutation steps, the worst (lowest fitness score) four chromosomes were replaced by the "elite" chromosomes.

Using the process described above, the genetic algorithm computes successive generations of chromosomes. Convergence is achieved when a chromosome (solution set) is found with a fitness function value of 1.0.

Air Traffic Data

Four hours of air traffic data were recorded from the Host Computer System (HCS) of the Memphis ARTCC (ZME) on 11 October 2000 from 1930 to 2330 UTC. Time coincident weather (wind and temperature) forecasts generated by the U.S. National Weather Service were also captured. The traffic data consisted of controller directives (e.g., flight plans, hold or interim altitude messages) and surveillance position reports of the aircraft (referred to as tracks). For this ZME recording, 1,749 flights were captured that had both track and at least one flight plan message.

Once this "raw" traffic data is captured from the field recording, it undergoes an extraction process. During the extraction process, as explained in [19] and [20], flight data is inserted into a relational database for each unique flight that has at least one flight plan message. If a flight entered ZME after the recording started, its first flight plan represents the first flight plan recorded by the ZME HCS. If a flight already had entered ZME when the recording started, its first flight plan represents a later flight plan amendment cleared by ZME controllers. The extraction ignores track messages preceding the first recorded flight plan; thus, many flights are truncated at the start of the scenario recording. Some flights are excluded altogether, which occurs if the flight had no flight plans recorded prior to their track.

This initial extraction process of the ZME data set yielded 1,694 flights. The analysis of this air traffic scenario required further culling. Four flights were removed due to significant errors in their track messages. The remaining set of 1,690 flights requires a defined start and end criteria for accurate analysis. The inbound hand-off time forms a consistent birth point for each flight, and the ZME boundary crossing point (or track end time for arrivals) forms a similarly consistent termination point. However, 14 of the flights were never under ZME control during the recording interval and were therefore excluded. Of the remaining flights, 62 never physically entered the ZME boundary and were therefore excluded as well.

The complete extraction process yielded 1,614 flights from the full ZME traffic recording; the corresponding recorded traffic data is called the extracted scenario. A time history of the aircraft count for the extracted scenario is shown in Fig. 2. The flight plan message and track position at hand-off time for each of the 1,614 extracted flights were

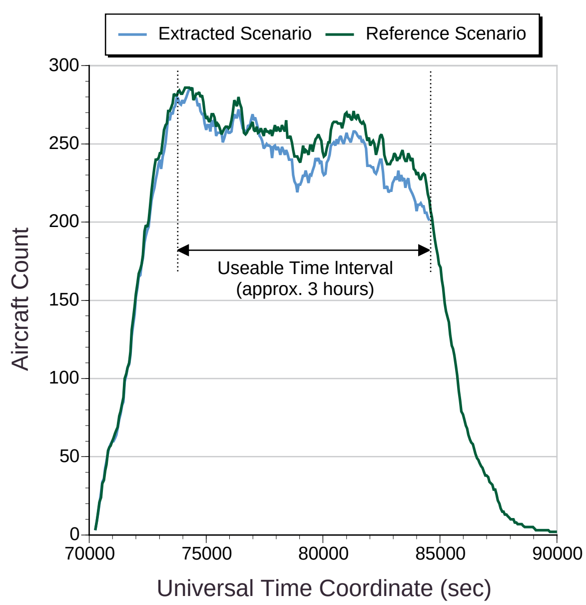

Figure 2. Aircraft Counts vs. Time

input into an air traffic simulation, along with the recorded weather forecasts. The resulting synthesized track data is called the reference scenario; a time history of the aircraft count for the reference scenario is shown in Fig. 2.

It is noted that the extraction process creates an artificial ramp up period in the traffic scenarios. This is a result of the requirement of a preceding flight plan before the track is captured. For both scenarios, the aircraft counts rise from practically zero to about 280 flights at 73,800 seconds in the recording; hence, any analysis should begin after this "steady state" time. Figure 2 also revealed that when the recording ended at about 84,600 seconds, the extracted scenario effectively ended, but the reference scenario continued its simulation of aircraft to either their ZME boundary crossing or landing within the ARTCC. The track data outside the start and end times specified above was excluded for the purpose of determining conflict properties of the reference scenario. Hence the four-hour ZME recording reduces to approximately three hours of useable traffic data containing 1,444 flights. The results of comparing the conflict properties of the reference and time-shifted conflict scenarios presented in the next section are based on these three hours of ZME traffic data.

Another observation is that Fig. 2 shows a modest bias of about ten aircraft in the reference scenario after approximately an hour after the steady state time (77,400 seconds). A likely explanation is that the wind forecasts used to produce the simulated tracks in the reference scenario have a proportional wind error after about two hours into the recording (i.e. 77,400 seconds). It is believed that this bias could be attenuated in the future with higher accuracy wind data, but is acceptable for this demonstration of the methodology.

The extracted scenario is input into the genetic algorithm, which computes a time shift or "delta time" for each flight to produce a set of conflicts with the desired properties (matching the reference scenario). These delta times are applied to the extracted scenario to create the time shifted scenario. Results are presented in the following section.

Results and Discussion

Figure 3 shows the distribution of time shifts generated by the genetic algorithm to match the primary properties of the reference conflict set. It found that almost half (47%) of the tracks were time shifted by less than 1 min, and that 99% of the tracks were time shifted by less than 5 min. The maximum time shift was under 7 min. These results indicate that close matching of primary properties can be accomplished with a minimal temporal perturbation of the recorded tracks.

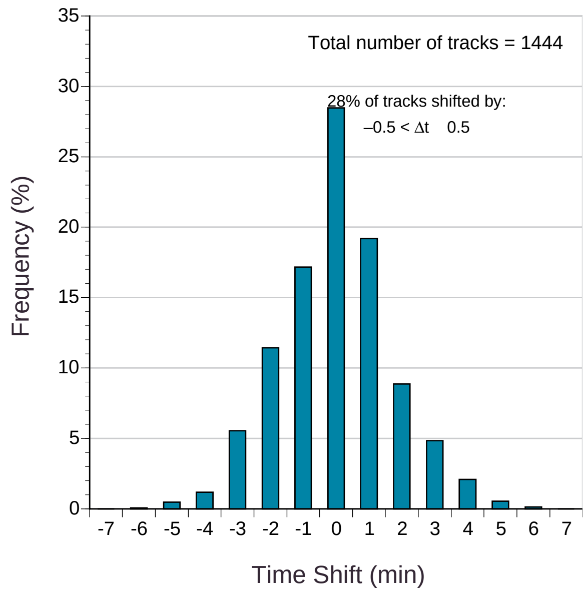

Figure 3. Histogram of Time Shifts

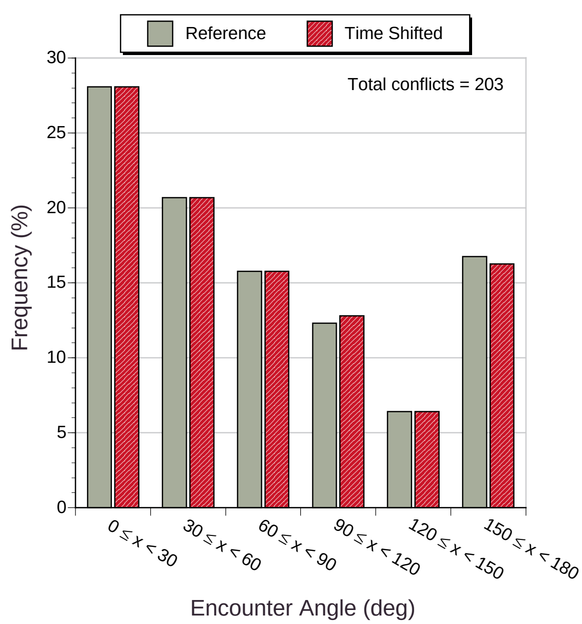

Figure 4. Encounter Angle Distributions

Primary Properties of the Conflict Sets

It is recalled the genetic algorithm determines the time shifts by attempting to match the primary property distributions (within a user-specified tolerance). A key primary property is the total number of conflicts. The reference set contained 203 conflicts, and it was found that the time shifted tracks also contained 203 conflicts (albeit not the same conflicts). Figures 4 to 7 present data on the distributions of the other primary conflict properties used in this work: encounter angle, minimum horizontal separation, minimum vertical separation, and vertical flight phase at conflict start. It is observed that the time shifted distributions match the reference distributions very well for all four properties.

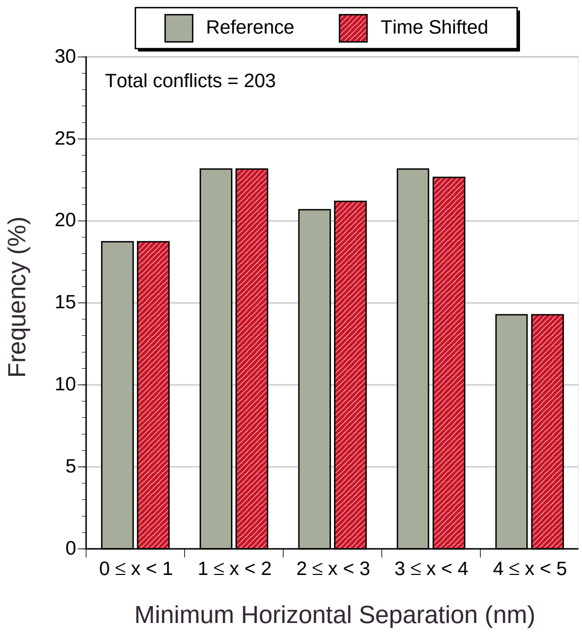

Figure 5. Min. Horiz. Separation Distributions

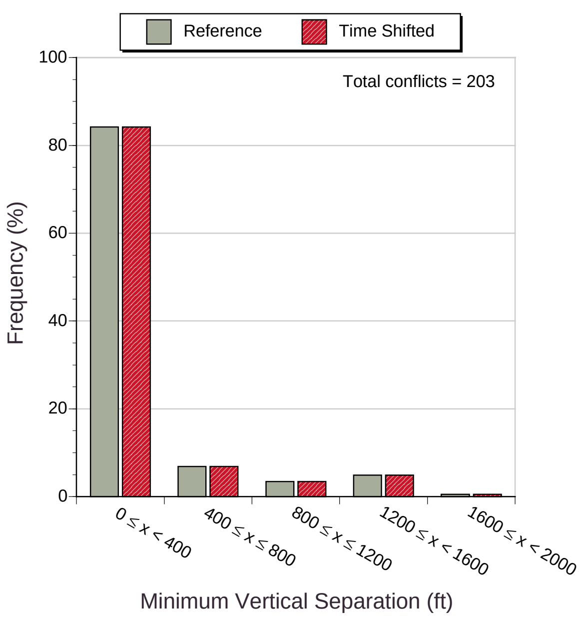

Figure 6. Min. Vert. Separation Distributions

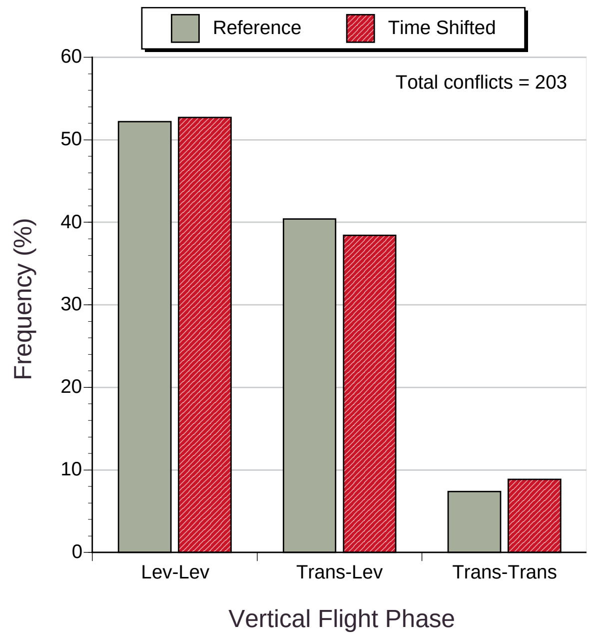

Figure 7. Vertical Flight Phase Distributions

Secondary Properties of the Conflict Sets

A conflict set has many properties other than the five primary properties identified in this work; they are considered to be secondary properties. The distinction between primary and secondary properties is that primary properties are those likely to have a major influence on the performance of any conflict probe, while the secondary properties are those likely to have a relatively minor influence on conflict probe performance.

Distributions of some secondary properties were determined and compared for the reference and time shifted sets. It is emphasized that the time shifting process made no attempt to match the secondary properties. The objective of this exercise is to see how well some secondary properties match up, as a by-product of the explicit matching process for primary properties. The secondary properties selected were: (1) total number of conflicting aircraft, (2) conflict duration, i.e., time interval of separation loss, (3) average horizontal position of conflict partners at conflict start, (4) average altitude of conflict partners at conflict start, and (5) conflict rate.

The total number of aircraft involved in conflicts in the reference set was 310, while the corresponding number in the time shifted set was 320 (which is only 3% off). Figure 8 presents data for conflict duration; it can be seen that there is a good match of the distributions. Figures 9 and 10 present data for average horizontal positions and altitudes,

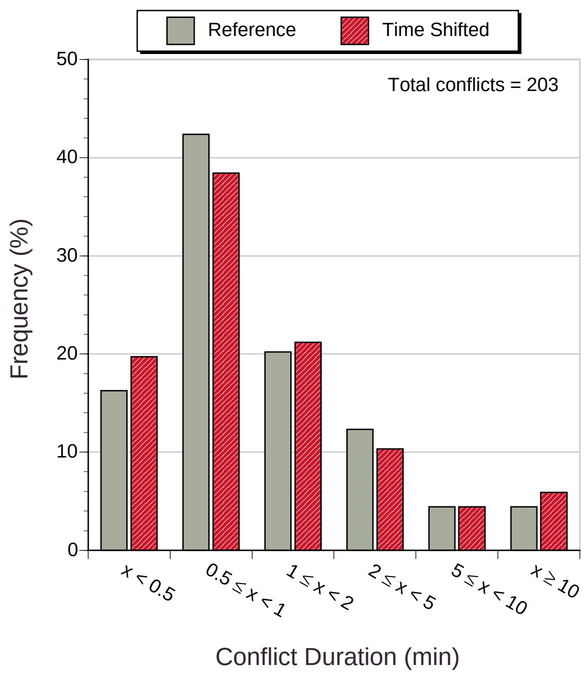

Figure 8. Conflict Duration Distributions

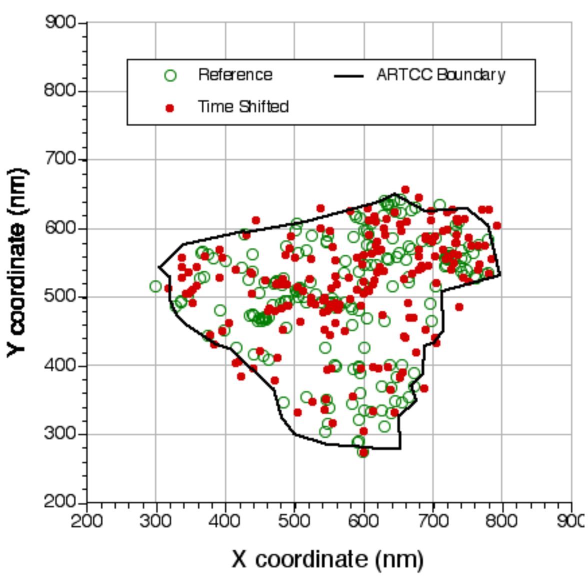

Figure 9. Conflict Position Distributions

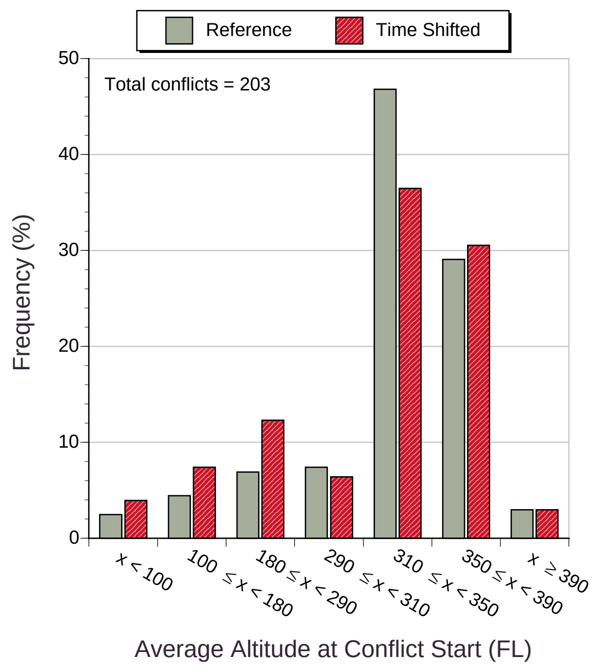

Figure 10. Conflict Altitude Distributions

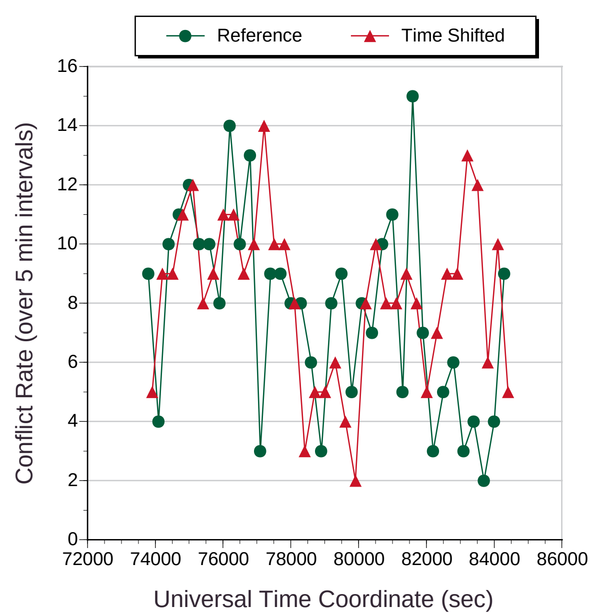

Figure 11. Conflict Rates vs. Time

respectively, at conflict start; there is a good qualitative match (of general trends) and a rough quantitative match (of actual values). Figure 11 presents data for conflict rate (number of conflicts over a 5 min interval); although the time variations do not match very well, it can be seen that the mean and range of the two data sets are very similar. The overall conclusion is that close matching of the chosen primary properties additionally provides a good matching of some secondary properties. This is further evidence that conflicts provided by the time shifted tracks reflect many of the essential characteristics of the reference conflict set.

Conclusions

A time shifting methodology has been developed for generating conflict scenarios using recorded air traffic data. The time shifting process was implemented by a genetic algorithm that attempted to match the primary properties of conflicts that would be observed in real air traffic operations if there were no controller actions to separate traffic. The primary properties used in this study were the number of conflicts, and the distributions of encounter angle, minimum horizontal separation, minimum vertical separation, and vertical flight phase at conflict start.

This methodology was successfully demonstrated using three hours of recorded air traffic data from the Memphis ARTCC. A key result of this work is that primary properties can be closely matched with very small time shifts in the track data. The demonstration study showed that 99% of the tracks were time shifted by less than 5 min. An interesting observation is that close matching of the primary properties used in this study additionally provides a good match for some other secondary properties. This indicates that the conflict set provided by the time shifted tracks reflects many of the essential characteristics of the reference conflict set, including some that were not explicitly matched.

References

[1] Kuchar, J.K. and L.C. Yang, 2000, "A Review of Conflict Detection and Resolution Modeling Methods," *IEEE Transactions on Intelligent Transportation Systems*, Vol. 1, No. 4, pp. 179-189.

[2] RTCA, 1995, *Free Flight Implementation*, Final Report of RTCA Task Force 3, Washington, D.C.

[3] Brudnicki, D.J. and A.L. McFarland, 1997, "User Request Evaluation Tool (URET) Conflict Probe Performance and Benefits Assessment," *1st USA/Europe Air Traffic Management R&D Seminar*, Saclay, France.

[4] McNally, B.D., R. Bach, and W. Chan, 1998, "Field Test of the Center-TRACON Automation System Conflict Prediction and Trial Planning Tool," Paper No. 98–4480, *AIAA Guidance, Navigation, and Control Conference*, Boston, MA.

[5] van Doorn, B., B. Bakker, and C. Meckiff, 2001, "Evaluation of Advanced Conflict Modelling in the Highly Interactive Problem Solver," *4th USA/Europe Air Traffic Management R&D Seminar*, Santa Fe, NM.

[6] Bilimoria, K.D., 2001, "Methodology for the Performance Evaluation of a Conflict Probe," *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 3, pp. 444-451.

[7] Loureiro, T., K. Blin, E. Hoffman, and K. Zeghal, 2001, "Development of a Tool for Comparing Conflict Detection Algorithms for Air Traffic Management," Paper No. 2001-4053, *AIAA Guidance, Navigation, and Control Conference*, Montréal, Canada.

[8] Cale, M.L., M.M. Paglione, H.F. Ryan, R.D. Oaks, and J.S. Summerill, 1998, "Application of Generic Metrics to Assess the Accuracy of Strategic Conflict Probes," *2nd USA/Europe Air Traffic Management R&D Seminar*, Orlando, FL.

[9] Alliott, J.-M., N. Durand, and G. Granger, 2001, "A Statistical Analysis of the Influence of Vertical and Ground Speed Errors on Conflict Probe," *3rd USA/Europe Air Traffic Management R&D Seminar*, Napoli, Italy.

[10] Mitchell, M., 1998, *An Introduction to Genetic Algorithms*, Cambridge, MA, MIT Press.

[11] Delahaye, D., J.-M. Alliot, M. Shoenauer, and J.-L. Farges, 1994, "Genetic Algorithms for Partitioning Air Space," *Proceedings of the 10th IEEE Conference on Artificial Intelligence for Applications*, San Antonio, TX.

[12] Giannazza, D., and J.-M. Alliot, 2002, "Optimization of Air Traffic Control Sector Configurations Using Tree Search Methods and Genetic Algorithms," *21st Digital Avionics Systems Conference: Air Traffic Management for Commercial and Military Systems*, Irvine, CA.

[13] Oussedik, S., and D. Delahaye, 1998, "Reduction of Air Traffic Congestion by Genetic Algorithms," *Fifth International Conference on Parallel Problem Solving from Nature*, Amsterdam, The Netherlands.

[14] Oussedik, S., D. Delahaye, and M. Schoenauer, 1998, "Air Traffic Management by Stochastic Optimization," *2nd USA/Europe Air Traffic Management R&D Seminar*, Orlando, FL.

[15] Brinton, C., J. Krozel, B. Capozzi, and S. Atkins, 2002, "Improved Taxi Prediction Algorithms for the Surface Management System," Paper No. 2002-4857, *AIAA Guidance, Navigation, and Control Conference*, Monterey, CA

[16] Garcia, J., A. Berlanga, J.M. Molina, J.A. Besada, and J. Casar, 2002, "Planning Techniques for Airport Ground Operations," *21st Digital Avionics Systems Conference: Air Traffic Management for Commercial and Military Systems*, Irvine, CA.

[17] Mondoloni, S. and S. Conway, 2001, "An Airborne Conflict Resolution Approach Using a Genetic Algorithm," Paper No. 2001-4054, *AIAA Guidance, Navigation, and Control Conference*, Montréal, Canada.

[18] Mondoloni, S., M.T. Palmer, and D.J. Wing, 2002, "Development of a Prototype Airborne Conflict Detection and Resolution Simulation Capability," Paper No. 2002-4446, *AIAA Guidance, Navigation, and Control Conference*, Monterey, CA.

[19] Oaks, R.D. and M.M. Paglione, 2002, "Generation of Realistic Air Traffic Scenarios Using a Genetic Algorithm," *21st Digital Avionics Systems Conference: Air Traffic Management for Commercial and Military Systems*, Irvine, CA.

[20] Oaks, R.D. and M.M. Paglione, 2001, "Generation of Realistic Air Traffic Scenarios Based on Recorded Field Data," *46th Annual Meeting of the Air Traffic Control Association*, Washington, D.C.

Key Words

Conflict probe evaluation; conflict scenario; genetic algorithm; recorded traffic data; time shifting.

Biographies

Mike M. Paglione is the Conflict Probe Assessment Team Lead in the Simulation and Analysis Branch (ACB-330) at the FAA William J. Hughes Technical Center. He has extensive experience in air traffic control automation algorithms, simulation problems, analysis of decision support software, applied statistics, and general systems engineering. He holds B.S. and M.S. degrees in Industrial and Systems Engineering from Rutgers University.

Robert D. Oaks is a Senior Computer Systems Analyst with Veridian IT Services, Inc., at the FAA William J. Hughes Technical Center. He has over thirty-five years of experience with all phases of software development most of which has been with real-time systems and the simulation of real time systems. He holds a B.A. degree in Mathematics from San Fernando Valley State College and a M.S. degree in Computer Science from the New Jersey Institute of Technology. He is a member of ACM, AIAA, and ATCA.

Karl D. Bilimoria is a Research Scientist at the NASA Ames Research Center, working on NASA's Airspace Systems program, where he has made significant research contributions in the area of Conflict Detection and Resolution. He holds a Ph.D. degree (1986) in Aerospace Engineering from Virginia Tech. From 1987 to 1994, Dr. Bilimoria was on the aerospace engineering faculty at Arizona State University, where he held the positions of Assistant Professor and Research Scientist, working in the areas of flight dynamics and optimal control. He is an Associate Fellow of the AIAA, and serves as an Associate Editor of the Journal of Guidance, Control, and Dynamics.